DOCUMENT RESUME

ED 277 997 CS 008 658

AUTHOR Bruce, Bertram
TITLE Deep Case Systems for Language Understanding. Technical Report No. 362.
INSTITUTION Bolt, Beranek and Newman, Inc., Cambridge, Mass.; Illinois Univ., Urbana. Center for the Study of Reading.
SPONS AGENCY National Inst. of Education (ED), Washington, DC.
PUB DATE Dec 85
CONTRACT 400-81-0030
NOTE 33p.
PUB TYPE Reports - Research/Technical (143)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Case (Grammar); *Cognitive Processes; Computer Uses in Education; *Language Processing; *Linguistic Theory; Reading Instruction; *Reading Research; *Structural Analysis (Linguistics); Theory Practice Relationship

ABSTRACT

Noting that the idea of linguistic case has evolved from an account of noun affixes (surface case) to an account of how syntactic relations between noun phrases and sentences map into deep relations between objects and events (deep case), this paper examines the notion of deep case as it applies to natural language understanding. Following a review of work on case grammars, primarily from the perspective of their use in computer natural language understanding systems, the paper discusses the distinction between surface cases and deep cases. It then examines the relation of deep cases to grammatical explanation and representation of meaning. The paper concludes with descriptions of four representational approaches to theories of case structures: (1) deep case relations, (2) case sequence paradigms, (3) discourse analysis, and (4) conceptual cases. Three figures and a four-page reference list are provided. (FL)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED277997

CENTER FOR THE STUDY OF READING

Technical Report No. 362

DEEP CASE SYSTEMS FOR LANGUAGE UNDERSTANDING

Bertram Bruce

Bolt Beranek and Newman Inc.

December 1985

University of Illinois
at Urbana-Champaign
51 Gerty Drive
Champaign, Illinois 61820

Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, Massachusetts 02238

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

This paper is based on an entry to appear in S. Shapiro (Ed.), Encyclopedia of Artificial Intelligence. New York: John Wiley & Sons, Inc. Publishers. I would like to thank Pat Chrosniak and Harry Blanchard for many useful comments on an earlier draft, and Tricia Neth for preparing the manuscript. The writing was supported in part by the National Institute of Education under Contract No. 400-81-0030.

CS008658

Abstract

This paper reviews work on case grammars, primarily from the perspective of their use in computer natural language understanding systems. A distinction between surface cases and deep cases in drawn. Deep cases are then discussed in relation to grammatical explanation and representation of meaning. Four representative approaches to theories of case structures are described.

Deep Case Systems

For Language Understanding

This paper examines the linguistic notion of case as it applies to natural language understanding. Case theory suggests an approach to the representation of sentence meaning and is important in accounting for the way the structure of sentences relates to those meanings. Most AI natural language systems make use of these ideas in some form (Bruce, 1975) and the concepts relate closely to the idea of schemata (Rumelhart, 1980) as applied in theories of reading. Applications of case theory have appeared in intelligent systems for such diverse areas as medical diagnosis (Chokhani, 1973; Kulikowski & Weiss, 1971) and speech understanding (Baranofsky, 1974; Nash-Webber, 1975). Research has addressed issues of efficiency, flexibility, scope, and grain (Bobrow & Winograd, 1977; Moore & Newell, 1973; Winograd, 1975).

The notion of "case" has been used to refer to several related concepts. Traditionally, it has meant the classification of nouns according to their syntactic role in a sentence, signalled by various inflected forms. In English, only pronouns have these case inflections. For instance, the first person singular pronoun is "I" (nominative case), "me" (accusative/objective case), or "my" (genitive/possessive case) according to its use as subject, object, or possessive article. In languages such as Greek, all nouns are given affixes which indicate their case.

The idea of a direct relationship between inflections and cases is one kind of case, also called "surface" or "syntactic level" case. For example, in the sentence

(1) Susan kicked the football with her foot.

each NP has a syntactic role:

- Subject: Susan
- Direct object: the football
- Object of the prepositon with: her foot

However, in understanding language, it is not sufficient to recognize the syntactic role of noun phrases (NPs). One would also like to know the semantic role each NP plays in the meaning of the sentence. Using the structural features together with lexical, morphological, and semantic information it is possible to determine that sentence (1) describes an event of kicking in which:

- Susan is the kicker, the agent;
- the football is the kickee, the object;
- her foot is used to perform the kicking, the instrument.

The latter analysis suggests another sense of "case" (also called "deep case," "semantic case," or "theta role"), namely, a categorization of noun phrases according to their conceptual roles in the action described by a sentence. Conceptual roles are independent of the particular verb or predicate being expressed. The agent case, then, is a generalization of many ideas: kicker, reader, walker, and dancer; i.e., one who

performs an action. Because deep cases describe meanings, rather than the words and structure that express those meanings, they are claimed to be language independent. Much of the discussion of deep cases has focused on identifying a small number of these conceptual roles which can be used for describing the meaning of any sentence. A set of deep cases is then called a "case system."

Surface Cases

One way to categorize nouns is by their endings or inflections. For the purposes of natural language processing, it is more useful to define surface case as a general syntactic categorization of noun phrases. Another way to think of surface case is as a property that is assigned to an NP, which is manifested in the sentence as a syntactic marker or signal, called a case marker.

Various linguistic elements can be case markers. The primary one is the case-affix, i.e., an ending attached to a noun form. Many would consider that prepositions (or postpositions) serve a similar function. Word order, as in English, can also be viewed as a case marker. In addition, case assignment interacts with such features as gender and definiteness of the noun phrase. This vew of case, then, generalizes the notion of surface case from simple noun inflections to a property which all NPs have and which may be expressed with word endings, word order, or other structural features.

How many distinct surface cases are there? One way to determine this is to consider a language in which cases are expressed by nominal inflections. In Latin, for example, five or six cases are usually distinguished: nominative, accusative, genitive, dative, ablative, and sometimes vocative.

But simply identifying surface cases is not that helpful in processing natural language, since surface cases are merely signals for which deep case to assign. In other words, for each conceptual role, one needs to account for the case markers that identify it. The degree to which a case based theory can account for inguistic behavior depends upon the way the cases mediate between surface forms and conceptual structures. The remaining discussion focusses on conceptual structures, or deep cases.

Deep Cases and Grammatical Explanation

The notion of deep cases is not new. For instance, Sonnencheins's demand that cases "denote categories of meaning" (see Jespersen, 1965) is in effect a statement that there are two levels of cases, the surface level indicated by case-affixes and a deeper level which may be common to more than one language. Fillmore (1968, pp. 2-3) presents a good argument for the universality of deep cases in natural language, saying that:

> What is needed is a conception of base structure in which case relationships are primitive terms of the theory and in which such concepts as "subject" and "direct object" are missing. The latter are regarded

as proper only to the surface structure of some (but possibly not all) languages.

Because deep cases focus on (conceptual) events rather than on syntactic constructions, they can help explain the relative "acceptability" of certain sentences. For example, one concept of the event "kicking," is that in (1). This concept encompasses such notions as agent, object, instrument, location, and so on. Knowledge of this concept, along with an understanding of concepts such as "football" and "foot," give an account of how to understand (1). At the same time it leads one to question sentences such as:

(2) Susan kicked the new idea.
(3) Susan kicked.
(4) Susan and her foot kicked the football.

Sentence (2) seems strange because the sense of "kick" used here seems to require a concrete object. Sentence (3) seems strange because the object of the kicking needs to be mentioned explicitly. Cues from the discourse as to what was kicked (or an intransitive interpretation of "kick") are needed to make the sentence comprehensible. Sentence (4) is also odd because, while either "Susan" or "her foot" or "Susan and Joe" could be the subject of the sentence, it is difficult to conjoin objects that play different roles in the meaning of the sentence.

These ideas can be formalized by postulating for each verb a case frame consisting of two elements:

- case structure. What are the case slots, or set of cases which play a role in the event denoted by the verb, e.g., a "kicking"? Which of these slots are optional, which are obligatory?
- selection restrictions. What are the semantic constraints on the objects which fill each slot in the case structure?

Selection restrictions may vary from global constraints on the use of a case with any predicate (e.g., "every agent must be animate") to local constraints on the use of a case with a particular predicate (e.g., "the object of 'spend' must be a resource"). Thus, for kicking, one might infer a case frame with the following slots and restrictions:

(5) [{agent}:animate object,
object:physical object,
{instrument}:physical object,
{source}:location,
{goal}:location]

The curly brackets are used here to indicate that the particular slot in the case structure is optional. This case frame says among other things, that if the agent of the kicking is described, it should be animate. Using an inanimate object description would suggest an interpretation in which that object is seen as animate. As discussed above, the prepositions and word order in a sentence may indicate which case is intended for each NP. If the indicated cases pass the appropriate selection restrictions and if they correspond to the cases allowed by the

case structure then the sentence should be easy to understand. Otherwise it can be considered ungrammatical or at least as grounds to re-interpret the event.

Most language understanding systems use case frames or equivalent mechanisms for semantic checking. A parser must check that the features of nominal constituents in the sentence satisfy the selection restrictions for the verb. The case frame may help to disambiguate among senses of the verb; either the case structure or the selection restrictions will distinguish the two senses. Furthermore, the selection restrictions can help the system identiry anaphoric references.

For instance, consider sentence (1). The indicated cases are [agent, object, instrument], each of which are present in (5) and the required object case is present. Susan is animate; the football and her foot are physical. Thus (1) can be easily mapped into the case structure for "kick."

Sentence (2) also indicates an acceptable case structure, [agent, object], but a new idea is not a physical object. Since the selection restrictions for the object slot are violated, the sentence is less easily mapped into the case frame and hence, less comprehensible. In contrast, sentence (3) obeys the selection restrictions of (5). Susan, the agent, is animate. However, its indicated case structure, [agent], does not contain the object case required by (5). Thus, it too, is problematic. For sentence (4), the case structure seems to be

[agent/instrument, object]. While either the agent or the instrument can be the subject of a kicking sentence, a case cannot be assigned to them when they are conjoined.

In general, discourse information can significantly alter the interpretation of a single sentence. If (2) were to follow a discussion of Susan's invention which wouldn't work, then the context might allow the "new idea" to be interpreted as a physical object. Or, suppose one describes Susan running towards a football, and then utters sentence (3). In that situation, one could easily infer that the object is the football.

Notice that in neither of these situations has the case structure or selection restrictions been violated, but, rather, the context provides information, which is missing from the sentence in isolation. Some language understanding systems allow ellipsis of the obligatory slots in case structures, i.e., if there is no filler, the system looks for nearby NPs to fill the slot.

Deep Cases and Meaning Representation

Underlying our discussion is the idea that people have a generic concept of an event such as kicking. Then, a sentence such as (1) serves to describe a particular instance of such an event. A formal representation of the concept of kicking can be given by first defining a predicate, kicking*, which represents the set of all events which are examples of kicking. The expression,

(6) (Ex) [kicking* (x)]

could then be read as, "there is an event which is an instance of kicking."

Usually, the kind of event is expressed as a verb, e.g., "to kick." However, an event description can also be realized as a noun phrase. One could say either "they prepared the meal" or "their preparation of the meal." By choosing events as primary entities, the semantic similarities among these phrases is captured naturally.

An event description points out an event and also distinguishes that event from other events of the same type by specifying various properties, or relationships between objects and the event. By asserting several propositions about the event, sentence (1), for example, indicates which kicking is being discussed. This set of propositions can be expressed as a conjunction of binary relations:

(7) (Ex) [kicking* (x)
& agent (x, Susan)
& object (x, the football)
& instrument (x, her foot)
& time (x, past)]

These relations suggest a formalism for representing sentence meaning. Some natural language understanding systems, assuming a small number of these fixed relations, parse sentences into their deep case structure rather than the traditional surface structure parse (see Figures 1 and 2).

Insert Figures 1 and 2 about here.

Because a class of verbs with related meanings can be used to describe similar events, these verbs share aspects of their case frames. For instance,

(8) Fred bought some pickles from Reuben.
(9) Reuben sold some pickles to Fred.

a case theory should capture the fact that sentences (8) and (9) describe the same event from a different perspective. The meaning of (8) could be represented as,

(10) (Ex) [exchange* (x)
& agent (x, Fred)
& goal (x, Fred)
& object (x, some pickles)
& source (x, Reuben)]

The meaning of (9) is similar, the only difference being that its agent is Reuben. Notice that this account requires that the subject have two deep cases. Jackendoff (1972, pp. 34-35), uses similar examples to justify his claim that an NP can have multiple deep cases.

Systems that make use of semantic similarities among verbs are described in (Hendrix, Thompson, & Slocum, 1973; Norman, Rumelhart, & the LNR Research Group, 1975). Identifying the case generalizations for classes of verbs based on cross linguistic evidence is the subject of ongoing research (Levin, et al., 1985).

One formalism for representing case frames is that of semantic networks. These were originally proposed by Quillian (1968) to capture the objective aspect of word meaning. The associative links between verb concepts (case frames) and real world knowledge facilitate inferences made from sentence meanings. Semantic network representations with structured inheritance, e.g., KL-ONE (Brachman, 1979), allow information about the syntactic and semantic regularities among verbs to be shared. Discussions of inferencing and case frame representation can be found in Charniak (1975) and Simmons (1972).

One problem is that an indefinite number of properties can be specified for a given event. For example,

(11) Because her arm hurt, Susan awkwardly kicked the football to Mary in the park rather than throw it.

could be represented,

(12) (Ex) [kicking* (x)
& reason (x, her arm hurt)
& agent (x, Susan)
& object (x, the football)
& time (x, past)
& manner (x, awkward)
& goal (x, Mary)
& location (x, the park)
& preference (x, throw it)]

Some of these properties distinguish one event from another while some merely modify, or provide additional information. For instance, the thing Susan kicks seems more significant than the fact that she kicks it awkwardly (her manner of kicking). Unfortunately, the labeling of a property as "distinguishing" or

"modifying" is rarely obvious. It is not difficult to imagine a context in which the manner in which an event happens is the distinguishing property and the object of the event is relatively insignificant. The distinction among properties is sensitive to the purpose of the speaker and the beliefs of both speaker and hearer. Nevertheless, there is often a strong intuition that certain properties belong with certain events. One could say that properties vary in their degree of binding to an event and that those properties which are most tightly bound are the deep cases.

Case Systems

Despite the compromises which seem necessary to dichotomize properties of events, there is a strong motivation to do so. By postulating a set of binary relations which represent the distinguishing properties of some generic event, one can define events as structures--known configurations which facilitate parsing and inference (Bruce, 1975; Martin, 1973; Norman, et al., 1975; Schank, 1974, 1973; Shapiro, 1971; Winograd, 1975). The complete set of deep cases available for describing events is called a case system. This section analyzes four significant case systems out of the many that have been proposed, in order to convey a sense of the range of work in this area. A more complete survey can be found in Bruce (1975).

Deep Case Relations

In classic papers, Fillmore (1968, 1971) has proposed a deep structure theory based on cases. A sentence in this deep structure consists of a modality plus a proposition:

(13) S -> M + P.

The modality constituent (M) includes negation, tense, mood, and aspect. The proposition (P) is a tenseless structure consisting of a verb and cases:

(14) $P \rightarrow V + C_1 + C_2 + \ldots + C_n$

where each C_i is a case name that generates either a noun phrase or an embedded S. There is a global constraint on rules of the form (14): At least one case must be present but no case may appear twice. Rules (13) and (14) are argued to be universal. Case markers are produced by the language specific Kasus element:

(15) $C_i \rightarrow K + NP$.

K generates a preposition, postposition, or case affix. One could generalize this notion to a Kasus function, which maps a deep structure proposition into a surface structure clause with possible word order changes.

Fillmore shows by example the deep case markers (Kasus functions) of various languages. He also gives some tentative rules for English. For example (1968, pp. 32-33):

> The A preposition is by; the I preposition is by if there is no A, otherwise it is with; the O and F [factitive case] prepositions are typically zero; the B [benefactive case] preposition is for; the D [dative case] preposition is typically to . . .

> If there is an A it becomes the subject; otherwise, if there is an I, it becomes the subject; otherwise, the subject is the O.

Fillmore makes an argument for deep case relations in analyzing verbs of any language, including English. He has proposed several systems which capture various aspects of the meaning of certain verbs. An example of his case systems appears in Table 1. In addition to these cases, there are also other relations "that identify the limits and extents in space and time that are required by verbs of motion, location, duration, etc." (Fillmore, 1971, p. 376).

Insert Table 1 about here.

Case Sequence Paradigms

Celce-Murcia (1972) developed a model for cases that emphasizes the fact that cluster of verbs appear to take similar sequences of cases. Her system is based on five deep case relations: causal-actant, theme, locus, source, and goal. Verbs are classified into paradigms according to the case sequences they allow. For example, the ergative paradigm consists of the sequences (for the active voice):

(causal-actant1, theme, causal-actant2)
(causal-actant1, theme)
(causal-actant2, theme)
(theme)

Note that a paradigm consists of both the case structure for the verb and constraints on the order of the case fillers. For

example, the ergative paradigm says that the theme can never precede the causal-actant.

"Break" is an example of an ergative verb. Thus,

(16) John broke the widow with a hammer.
(17) John broke the window.
(18) The hammer broke the window.
(19) The window broke.

are all well-formed since in each sentence one of the case sequences is matched (where "John" is the causal-actant1, "window" is the theme, and "hammer" is the causal-actant2).

Another example is the reflexive-deletion paradigm, in which the theme is deleted if it matches the causal-actant1. Thus "run" may be used in several ways:

(20) John ran to school.
(21) John ran a machine.
(22) The machine ran.
(23) The brook ran.

In each of the sentences there is a theme--John, machine, or brook. The paradigm allows the deletion of the theme if it is the same as the causal-actant. Thus the paradigm is

(causal-actant, goal)
(causal-actant, theme)
(theme).

Discourse Analysis

Grimes (1972) has developed a case system to serve as a foundation for discourse analysis. The definitions of the cases and their organization reflect his concern with event and episode representations. Grimes distinguishes between two kinds of generic events each with its own set of roles or deep cases.

Motion/position events have orientation roles and changes of state have process roles. In addition, the agent and benefactive roles are common to all events. These cases are shown in Table 2.

Insert Table 2 about here.

The following examples illustrate the use of these cases:

(24) The letter (O) fell to the floor (G).
(25) His house (O) is situated on top of a hill (R).
(26) The tide (V) floated the oil slick (O) into the harbor (G).
(27) This idea (O) came to me (G) from Austin Hale (S).
(28) This book (P) costs three dollars (Rf).
(29) She (A) makes dresses (P Re) from flour sacks (P M).
(30) Fred (A) fixed the engine (P) with this screwdriver (I).
(31) Sally (A) handed John (G) the biscuits (O).
(32) He (A) parted the rope (P G) with an axe (O I).
(33) The girl (P) died of malaria (F).
(34) The milk (P) turned sour on me (B).
(35) We (A) talked about politics (Rf).
(36) A breeze (O) came to him (G) from the sea (R).

The cases Grimes distinguishes are strongly influenced by linguistic, not conceptual, considerations, e.g., in (27) the transfer of the idea is not a physical movement. Sentence (27) has the same surface form as (36), which is a description of a physical transfer, so the two have similar case assignments.

Grimes also suggests the possibility of a more tightly defined role structure based on certain similarities in the roles: "The roles set up for orientation all have counterparts on the process side, and vice versa. Both kinds could be considered complementary variants of a single set of roles . . . Object and patient both identify what is affected, the one in

terms of motion or position and the other in terms of change of state in a process" (Grimes, 1972). These observations suggest the combined role structure shown in Figure 3.

Insert Figure 3 about here.

Conceptual Cases

Schank's (1971, 1973, 1974, 1975) cases, unlike those of Fillmore (1968) or Celce-Murcia (1972) are purely conceptual. Neither the primitive act nor its cases need be explicitly mentioned in an utterance. Instead, the argument for conceptual cases depends upon considerations of the pragmatics of human communication. One postulates a conceptual case because it is a relation relevant to the typical kinds of tasks which people address via language.

An essential element of most communication is the description of actions. Our knowledge of actions implies a "conceptual structure" built out of actions and their role fillers:

> ACTORS perform ACTIONS
> ACTIONS have OBJECTS
> ACTIONS have INSTRUMENTS
> ACTIONS may have RECIPIENTS
> ACTIONS may have DIRECTIONS. (Schank, 1974, p. 6)

One kind of conceptual structure or "conceptualization" comprises an act, with its "actor," and the relations "object," "direction," and either "recipient," or "instrument." Each of

these relations must be present (except that only one of direction or recipient is present).

Schank argues that a small number of concepts corresponding to "primitive acts" can be used to construct meaning representations for most descriptions of events. These primitive concepts are simple actions of the kind "move a body part" (MOVE), "build a thought" (MBUILD), "transfer a physical object" (PTRANS), and "transfer mental information" (MTRANS). The primitive ACTS together with the conceptual cases are the components of meaning representation with a "unique representation" feature: "We have required of our representation that if two sentences, whether in the same or different language are agreed to have the same meaning, they must have identical representations" (Schank, 1974, p. 4). It is questionable whether such a criterion can be met non-trivially. Do distinct utterances (by different speakers using different phrasings, at different times, in different situations) share significant portions of a conceptual network? Furthermore, a non-redundant representation such as Schank's raises serious questions of both psychological validity and efficiency for diverse tasks. Nevertheless, in many cases the mapping of utterances to conceptualizations seems to be exactly the process which humans exhibit. The unique representation also facilitates general inferencing by reducing the number of cases to be considered:

> The use of such primitives severely reduces the inferences problem in Al . . . since inference rules need only be written once for any ACT rather than many times for each verb that references that ACT. For example, one rule is that if you MTRANS something to your LTM [long term memory], then it is present there (i.e., you know it). This is true whether the verb of MTRANSing was see, hear, inform, remember or whatever. The inference comes from the ACT rather than the verb. (Schank, 1974 p. 10)

Conclusion

The notion of case has evolved from an account of noun affixes to an account of how syntactic relations between NPs and sentences map into deep relations between objects and events. These ideas have been applied to natural language processing for semantic checking and meaning representation.

In many language systems, a case frame is associated with each verb (and sometimes nouns). In recognizing the syntactic role of an NP in a sentence, the parser uses the case frame to verify that the semantic properties of the NP are consistent with some case which can occur in that syntactic position. This process can be used to block a parse path, to reject a sentence as ungrammatical, and to identify constraints for an ellipsed item or the referent of a pronoun.

Deep case systems are also an attempt to identify a fixed number of conceptual roles which can be used to describe any event. Representing deep cases as binary relations thus provides a formalization of the meaning of a sentence. This structure for describing knowledge has led to extensive research on semantic networks for knowledge representation. Out of this research, standard techniques for parsing and understanding have evolved to the extent that most current natural language systems incorporate these techniques in some form.

An important aspect of any case system is an account of how the deep cases are realized in a sentence. Many issues related to this accounting remain unresolved, such as whether an NP can have multiple cases (Jackendoff, 1972), and how to capture the regularities in the way the cases are realized (Levin, 1985). Although case grammar per se is not a major focus of current research within artificial intelligence, issues such as these are actively pursued in related work on schema theory, relational grammar, lexical functional grammar, generalized phrase structure grammar, and semantic grammar.

References

Baranofsky, S. (1974). Semantic and pragmatic processing in the SRI speech understanding system (Tech. Rep.). Stanford, CA: Stanford Research Institute.

Bobrow, D. G., & Winograd, T. (1977). An overview of KRL, a knowledge representation language. Cognitive Science, 1, 3-46.

Brachman, R. J. (1979). On the epistemological status of semantic networks. In N. V. Findler (Ed.), Associative networks. New York: Academic Press.

Bruce, B. (1973). Case structure systems. In Proc. IJCAI-73 (pp. 364-371). Stanford, CA: International Joint Conference on Artificial Intelligence, Inc.

Bruce, B. (1975). Case systems for natural language. Artificial Intelligence, 6, 327-360.

Celce-Murcia, M. (1972). Paradigms for sentence recognition (Tech. Rep. HRT-15092/7907). Santa Monica, CA: System Development Corporation.

Charniak, E. (1975). A brief on case (Tech. Rep. No. 22). Switzerland: Institute per gli studi Semantici e Cognitivi.

Chokhani, S. (1973). The definition and interpretation of a causal network model for disease within the CHRONOS system (Tech. Rep. NIH CBM-TM-17). New Brunswick, NJ: Rutgers University, Computer Science Department.

Fillmore, C. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), Universals in linguistic theory (pp. 1-88). New York: Holt, Rinehart & Winston.

Fillmore, C. (1971). Types of lexical information. In D. D. Steinberg & L. A. Jakobovits (Eds.), Semantics: An interdisciplinary reader (pp. 370-392). London: Cambridge University Press.

Grimes, J. (1972). The thread of discourse (Tech. Rep. NSF 1). Ithaca, NY: Cornell University.

Hendrix, G., Thompson, C., & Slocum, J. (1973). Language processing via canonical verbs and semantic models. In Proc. IJCAI-73, (pp. 262-269). Stanford, CA: International Joint Conference on Artificial Intelligence.

Jackendoff, R. S. (1972). Semantic interpretation in generative grammar. The MIT Press.

Jespersen, O. (1965). The philosophy of grammar. New York: Norton.

Kulikowski, C., & Weiss, S. (1971). Computer based models for glaucoma (Tech. Rep. CBM-TR-3). New Brunswick NJ: Rutgers University, Computer Science Department.

Martin, W. A. (1973). Translation of English into MAPL using Winograd's syntax, state transtition networks, and a semantic case grammar (Tech. Rep. Internal Memo 11). Cambridge, MA: MIT., Project MAC, Automatic Programming Group.

Moore, J., & Newell, A. (1973). How can MERLIN understand? In L. Gregg (Ed.), Knowledge and cognition. Los Angeles, CA: Lawrence.

Nash-Webber, B. (1975). Semantic support for a speech understanding system. In D. G. Bobrow & A. Collins (Eds.), Representation and understanding: Studies in cognitive science (pp. 351-362). New York: Academic Press.

Norman, D. A., Rumelhart, D. E., & the LNR Research Group (1975). Explorations in cognition. San Francisco, CA: Freeman.

Quillian, M. R. (1968). Semantic memory (Tech. Rep. AFCRL-66-189). Cambridge, MA: Bolt Beranek & Newman.

Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), Theoretical issues in reading comprehension (pp. 33-58). Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, R. (1971). Finding the conceptual and intention in an utterance in natural language conversation. In Proc. IJCAI-71, (pp. 444-454). London: International Joint Conference on Artificial Intelligence, Inc.

Schank, R. (1973). The fourteen primitive actions and their inferences (Tech. Rep. AIM-183). Stanford, CA: Stanford University.

Schank, R. (1974). Causality and reasoning (Technical Report No. 1). Castagnola, Switzerland: Instituto per gli studi Semantici e Cognitivi.

Schank, R. (1975). Conceptual information processing. New York: North Holland.

Schank, R., Goldman, N., Rieger, C. III, & Riesbeck, C. (1973). MARGIE: Memory analysis, response generation, and inference on English. In Proc. IJCAI-73, (pp. 255-261). Stanford, CA: International Joint Conference on Artificial Intelligence, Inc.

Schapiro, S. (1971). A net structure for semantic information storage, deduction and retrieval. In Proc. IJCAI-71, (pp. 512-523). London: International Joint Conference on Artificial Intelligence, Inc.

Simmons, R. F. (1972). Semantic networks: Their computation and use for understanding English sentences (Tech. Rep. CAI NL-6). Austin, TX: University of Texas, Computer Science Department.

Winograd, T. (1975). Frame representations and the declarative/procedural controversy. In D. G. Bobrow & A. Collins (Eds.), Representation and Understanding: Studies in Cognitive Science (pp. 185-210). New York: Academic Press.

Table 1: Fillmore's (1971) case system

Agent(A)	the instigator of the event
Counter-Agent(C)	the force or resistance against which the action is carried out
Object(O)	the entity that moves or changes or whose position or existence is in consideration
Result(R)	the entity that comes into existence as a result of the action
Instrument(I)	the stimulus or immediate physical cause of an event
Source(S)	the place from which something moves
Goal(G)	the place to which something moves
Experience(E)	the entity which receives or accepts or experiences or undergoes the effect of an action

Table 2: Grime's (1972) case system

Orientation Roles:	
Object(O)	the thing whose position or motion is being described
Source(S)	the location of the object at the beginning of a motion
Goal(G)	the location of the object at the end of a motion
Range(R)	the path or area traversed during a motion
Vehicle(V)	the thing which conveys the object and moves along with it
Process Roles:	
Patient(P)	the thing changed by a process or the thing whose state is being described
Material(M)	the thing changed by a process in its state before the change
Result(Re)	the thing changed by a process in its state after the change
Referent(Rf)	the field or object which defines the limitation of a process (as opposed to the thing affected by the process)
The Agentive Complex:	
Agent(A)	the one who is responsible for an action
Instrument(I)	the tool used in performing an action
Force(F)	the noninstigative cause of an action
The Benefactive Role:	
Benefactive(B)	the someone or something on whom an action has a secondary effect

Figure Captions

Figure 1. A traditional phrase structure parse for "Susan kicked the football with her foot."

Figure 2. A case oriented parse for "Susan kicked the football with her foot".

Figure 3. Interrelationships among roles (Grimes, 1972).

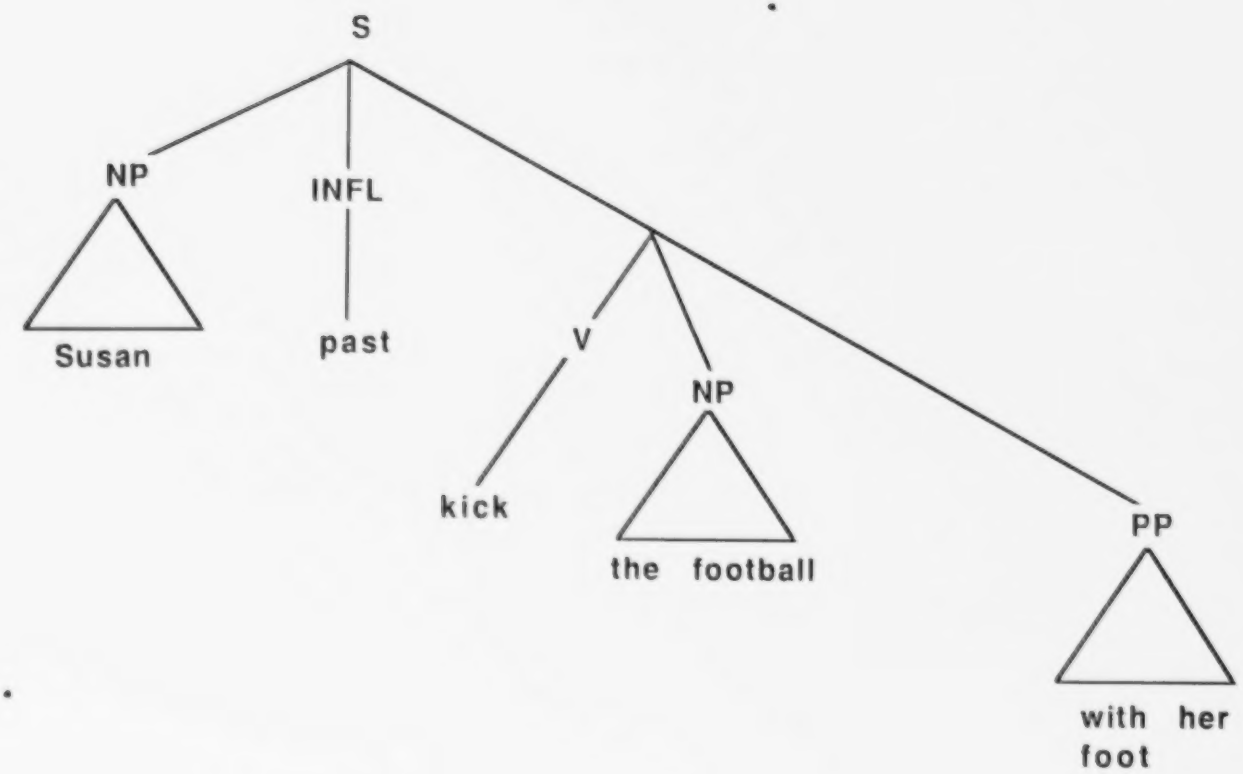
S
NP
INFL
Susan
past
V
NP
kick
the football
PP
with her foot

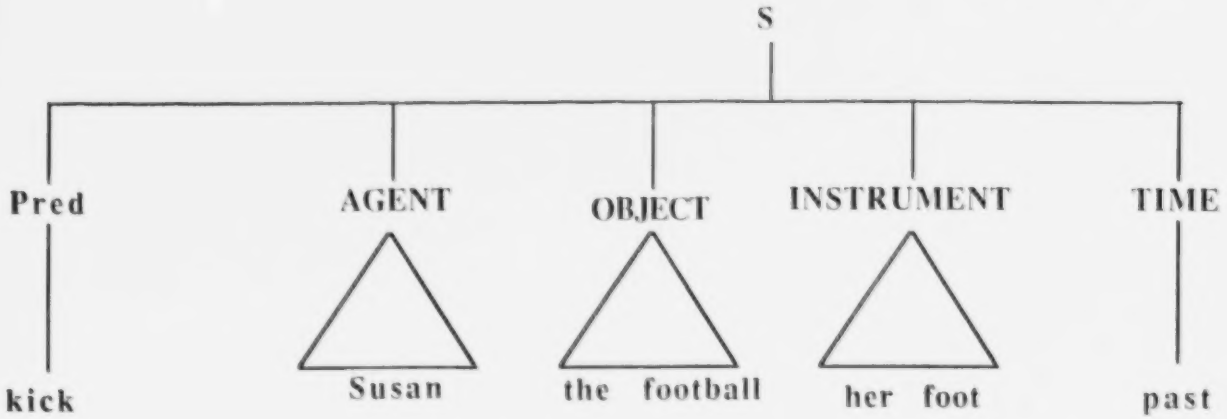
S
Pred
AGENT
OBJECT
INSTRUMENT
TIME
kick
Susan
the football
her foot
past

Orientation		Combined		Process
		A agent |		
		| Fc force		
		I instrument |		
V vehicle	-->	V vehicle		
O object	-->	P patient	<--	P patient
S source	-->	F former	<--	M material
G goal	-->	L latter	<--	Rs result
R Range	-->	R range	<--	Rf referent
		B benefactive		